DOCUMENT RESUME

ED 328 216 IR 014 811

AUTHOR Baker, Michael
TITLE Arguing with the Tutor: A Model for Tutorial Dialogue in Uncertain Knowledge Domains. Centre for Information Technology in Education Report No. 51.
INSTITUTION Open Univ., Walton, Bletchley, Bucks (England). Inst. of Educational Technology.
PUB DATE Sep 88
NOTE 57p.; For related reports, see IR 014 806-808 and IR 014 812-813.
PUB TYPE Reports - Research/Technical (143)

EDRS PRICE MF01/PC03 Plus Postage.
DESCRIPTORS *Computer Assisted Instruction; *Computer System Design; *Dialogs (Language); Discovery Learning; Expert Systems; Foreign Countries; Learning Theories; *Man Machine Systems; *Models; *Programed Tutoring
IDENTIFIERS *Intelligent Tutoring Systems

ABSTRACT

This report describes a specific approach to two current trends in Intelligent Tutoring Systems (ITS) research: (1) an emphasis on metacognitive skills which are not domain specific; and (2) a move away from viewing ITS research as being exclusively concerned with "knowledge communication." A human-computer dialogue model designed to teach the metacognitive skill of critical argument is described. The dialogue model occurs in a knowledge domain consisting of a set of justified beliefs, and is exemplified in the domain of teaching musical structures. The model is based on a dialogue planning system that involves higher level planning decisions, making claims and challenges, and understanding the components of dialogue such as actions, effects, and preconditions. Integrating intentional and topical dialogue structures, the model may also be used with a variety of tutorial styles within a guided-discovery tutoring approach. (39 references) (Author/DB)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

CENTRE FOR INFORMATION TECHNOLOGY IN EDUCATION

ED328216

IET

Arguing with the tutor: a model for tutorial dialogue in uncertain knowledge domains

by

Michael Baker

CITE Report No. 51

Institute of Educational Technology
The Open University
Milton Keynes
MK7 6AA
Tel: 0908 74066

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY
E.S. Henderson
TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

DRAFT DOCUMENT

Arguing with the tutor: a model for tutorial dialogue in uncertain knowledge domains

Michael Baker
Centre for Information Technology in Education
I.E.T.
The Open University
Milton Keynes MK7 6AA
tel. (0908) 653488
mj_baker@uk.ac.open.acs.vax

September 1988

This is a draft document for a chapter to appear in Guided-Discovery Tutoring. (ed.) M. Elsom-Cook. The author would therefore greatly appreciate any comments at this stage.

DRAFT DOCUMENT

Abstract

This chapter describes a specific approach to following two current trends in Intelligent Tutoring Systems research, which are (1) an emphasis on metacognitive skills which are not domain specific, and (2) a move away from viewing ITS research as being exclusively concerned with "knowledge communication" (Wenger 1987).

A human-computer dialogue model is described, designed to teach the metacognitive skill of critical argument, in a knowledge domain consisting of a set of consistent justified beliefs (exemplified in in the domain of teaching musical structures). The model is based on a dialogue planning system, which integrates intentional and topical dialogue structures, and which can flexibly integrate a variety of tutorial styles within a Guided-Discovery Tutoring approach.

Arguing with the Tutor:
a model for tutorial dialogue in uncertain knowledge domains

Michael Baker, September 1988

Contents

Introduction: dialogue and domain models 1
- **Dialogue models and Guided Discovery Tutoring** 1
- **Domain models and dialogue models** 6
- **An example domain: knowledge of musical structures** 8

A model for generating human-computer critical arguments 10
- **Remarks on methodology** 11
- **Representing the domain within the dialogue model** 11
- **The dialogue model** 15
- **Dialogue goals for critical argument** 16
  - *Higher level planning decisions* 18
  - *Making claims* 19
  - *Making challenges* 20
- *Components of dialogue goal operators* 22
- (i) actions 22
- (ii) Effects 22
- (iii) preconditions 23
- **An example hand-simulated dialogue** 27
- **The algorithm for dialogue generation** 29

A critique of the model in relation to existing work 33
- **Topic and goal-based approaches to dialogue** 33
- **Integrating the two approaches** 34
- **Symmetry, cooperativity and negotiation** 36
- **Symmetry and belief revision** 37
- **The succession of dialogue units** 38
- **Conversants' expectations and adjacency pairs** 40
- **Tutorial strategies, dialogue goals and related work** 41

Conclusions 43

Acknowledgements 46

References 46

Arguing with the Tutor:
a model for tutorial dialogue in uncertain knowledge domains

Michael Baker, September 1988

Introduction: dialogue and domain models

Dialogue models and Guided Discovery Tutoring

In guided discovery tutoring, the medium of interaction is the *dialogue model*. Within a computer-based environment, the dialogue model consists of an overall *structure* to the guidance provided to the student, and a set of *mechanisms* or algorithms for generating and sustaining dialogue. In comparison with the kinds of practical educational activities in which teachers and learners may engage in 'real life' settings, the dialogue model is used to provide guidance for a student's educational exploration within activitities which are themselves computer-based. The problem of providing guidance is therefore usually understood to be that of facilitating a student's educational exploration of a knowledge domain represented in the form of a computer-based *graphical interface*. In this chapter we shall base our discussion around the specific problem of providing guidance on the use of a *graphic music-editing tool*, one of whose roles is to allow the student to indicate musical structures in a melody, and to hear the melody played via a synthesiser linked to the system. The graphic inteface is thus the means by which the student can explore or 'directly manipulate' (Shneiderman 1982) the knowledge domain, and the dialogue model is used to give guidance on such exploration.

The fact that the dialogue model is designed for structuring human-*computer* interaction (as opposed to human-human interaction) imposes special constraints on the form of the model. Let us first examine what we mean by a dialogue model, before discussing constraints on such models within computer-based environments. The dialogue model is not (necessarily) a model for generation of natural language utterances, but refers to a much more general, abstract and cognitively based conception of the structure of human-computer interaction. Typically, such a model may rely on text formats to generate linguistic utterances, or even integrate illustrative examples generated from the knowledge domain itself. Dialogue structures occur on many levels of scale. At one level, there is the micro-structure of the succession of *speaker turns*, possibly with a single speaker taking several

simultaneous turns. At intermediate levels, structures can be identified in terms of the *goal*, purpose or intention of a succession of sentence-level utterances, and in terms of pursuit of successive *topics* . At the largest level (for tutorial dialogues), we can characterise *tutorial strategies* in terms of the consistent choice of one or more types of dialogue goal within topic constraints. For example, a tutorial strategy 'Socratic tutoring' might generate a large scale dialogue structure from consistent choice of the dialogue goals of eliciting and critiquing utterances made by the student on a specific topic.

The claim that tutorial guidance should be effected via a principled and structured interaction, considered in abstraction from mere linguistic expression of the subject domain, may seem thus obvious; but it is important to realise that most existing Intelligent Tutoring Systems (ITS) lack such a principled model of the interaction structure. Instead, they have concentrated on developing systems with a single tutorial style - such as 'Socratic tutoring' (1983) 'critiquing' (Miller 1984) and 'successive refinement' (Bonar & Cunningham 1988), or else the interaction is essentially structured by traversal of the domain knowledge representation. This single tutorial style is also usually expressed in terms of a succession of utterances of a specific kind, which lack a more general model of the succession of utterances, or *dialogue coherence*. For example, Socratic tutoring, as exemplifed in Collins' work (1977), consists essentially of generating a succesion of *questions*, with a larger scale structure which is determined by traversal of remaining topics to be taught in the (semantic network) domain knowledge representation. Their model of tutorial strategies involves deciding what *kind* of questions to ask, and on what topic - consider their "entrapment" strategies, for example (Collins 1977). There is no explicit representation of dialogue focus, and no attempt is made to accomodate the student's pursuit of learning goals within the dialogue.

We would argue that the structure of the tutorial interaction must be considered *sui generis* as based on a much more general model of human-computer interaction. The model should aim to fulfill (at least) the following requirements:

(1) the structure of tutorial dialogues should not simply aim to impart the subject

domain (by traversal of the domain knowledge representation);

(2) a dialogue model should be *generalisable* to other subject domains which share common cognitive characteristics;

(3) the dialogue model should attempt to synthesise a number of levels and types of structures;

(4) the dialogue model should aim to support a variety of possible tutorial styles and learning styles;

(5) the dialogue model should aim to be *symmetrical*, in terms of allowing the student to pursue the same set of dialogue goals (to achieve learning goals) that the system may employ (to achieve tutorial goals).

We shall return to each of these points many times in the ensuing discussion, but let us now briefly elaborate each of these claims in turn.

(1) A number of approaches to generation tutorial dialogues have been essentially 'bottom up', in that their structure derives directly from traversal of the domain knowledge representation. For example, in the WHY system Stevens and Collins (1977) produced a dialogue structure which was essentially based on traversing a series of scripts and subscripts relating to causal factors involved in rainfall. We argued earlier that such a topic-based level of structure is in fact an essential component of dialogue, and Collins & Stevens work here was in fact supported by protocol analysis. However, this single level of structure is not sufficient: we need to do justice to the student's and system's required intentional structures, and to support a variety of utterance types and tutorial styles.

(2) The claim that a dialogue model should be *generalisable* to other subject domains is both a statement of what we mean by a dialogue model here, and a statement of preferred methodology in ITS research. Our working hypothesis is that the dialogue model must be based on a *cognitively realistic* model of human dialogue, which follows trivially from the intuitively acceptable premise that a computational dialogue must be intelligible to the human participant. We would argue that, just as for a human dialogue participant dialogue

generation and understanding depends on accessing knowledge residing in long-term memory structures, a humanly intelligible computational dialogue should be generated from a cognitively realistic knowledge representation. (By the term 'cognitively realistic' we simply mean that the model should be at least a reasonable candidate cognitive model). In other words, given a cognitively realisitic dialogue model, and a certain relationship between thought and its linguistic expresssion, an intelligible dialogue structure would be based on a cognitively realistic model of the *subject domain*. Now, to continue this line of argument, given that we accept the premises

- computational knowledge representations should be based on cognitive models of human memory
- a computational dialogue must be intelligible to a human participant
- there is a close (complex) relationship between thought and its linguistic expression (*re* knowledge representation and dialogue structure)
- there are generalisable cognitive structures in memory and in language production and understanding, across knowledge 'domains'

then it follows that

computational dialogue structures and mechanisms should be generalisable across knowledge domains

In other words, a dialogue model should be generalisable. From the existing literature on human-computer interaction, there may be a number of possible claims with which the above claim may be confused. Our claim is *not* equivalent to the following claims:

(i) human-*computer* dialogue structures must be equivalent to human-*human* dialogue structures;

(ii) a different dialogue model is required for each domain.

With respect to (i), we agree that there are a number of special features of human-computer dialogue, which are not shared with human-human dialogues (as Suchman 1985, would argue, for example). For example, human-computer dialogues are unlike human conversations, in that they are written, and do not occur in real-time. We are making the much more general claim, that whatever these differences, the dialogue still needs a structure

which is intelligible to the human participant, and that this can best be achieved by maintaining a close relationship between cognitively realistic domain and dialogue models.

Considering the second possible misinterpretation (ii), we would argue that existing theories of human memory (such as semantic network models, Quillan 1968) are advanced as applying across a number of possible knowledge domains, and that therefore, the attendant dialogue model should also do so.

We mentioned that there was a methodological issue at stake here, in the preferred generalisability of dialogue models. Ohlsson (1986) has argued convincingly that we should be trying to produce tutoring systems which are generalisable, as the best way for the field to advance as a science. Clearly, what is the case for the whole system should also be so for its components, including domain, dialogue and student models. Another way of characterising this generalisability issue is in terms of theories of (human) *learning*, which should be embodied in any ITS: a general theory of human knowledge acquisition must depend on general theories of communication and representation of knowledge. In addition, we can also find support for our claims from the literature on education, which suggests that human teachers use the same teaching strategies (which we subsume within the dialogue model) in different domains (Bellack, Kliebard, Hyman & Smith 1966).

In summarising this point, it may be remarked that the existing ITS literature contains many examples of systems which have failed to promote learning largely because the dialogue mechanisms which have been alleged to facilitate such learning have been based on domain representations which were not *communicable* (Wenger 1987) or cognitively realistic (the use of production rules in the earlier versions of Clancey's 1987 GUIDON is a case in point).

(3) That dialogue models should attempt to synthesise a number of levels of structure (tutorial styles, intentional and topical) follows from the earlier arguments that the model should be cognitively realistic. The model should facilitate a student's pursuit of learning goals via

intentional structures, and should be perceived as coherent in terms of topical structures.

(4) In elaborating claim (2) we have already discussed reasons why the dialogue model should aim to support a variety of possible tutorial styles and learning styles. The claim follows from the generalisability of dialogue models, is supported by evidence that teachers use a variety of tutorial styles across domains, and is required to facilitate pursuit of alternative learning styles by the student.

(5) Finally, we would claim that the concept of *symmetry* within a dialogue model is an ideal which should be pursued, but one which has not been completely achieved in this research. By symmetry, we mean that the student should be able to pursue the same set of dialogue goals (to achieve learning goals) that the system may employ (to achieve tutorial goals), and additionally, that the system must be prepared to modify its beliefs in the way in which the student is expected to modify beliefs in response to teaching interactions. In the concluding section to this paper we will discuss how existing tutoring sytems lack such symmetry, and how the ideal is not simply an ideal, but is required for effective intelligent tutoring

Domain models and dialogue models

In this chapter we shall explore the relationship between domain and dialogue models in a specific domain - *knowledge of musical structure* - and describe a dialogue model which exploits specific properties of a cognitively realistic representation of that domain. The specific characteristics of this domain are that the knowledge of musical structures which human listeners and analysts possess is inherently *uncertain* and *ambiguous*. For example in the following well known musical extract (Figure 1), human analysts are undecided as to which of the two structural analyses (marked a and b) is preferred:

Figure 1

Mozart piano sonata K331, A major,
first 3 bars

In other words, it would be cognitively *un*realistic to have a domain representation for musical structure in an ITS which consisted of a single set of declarative 'facts' concerning the structure of each piece of music, which were assumed to be certain and unambiguous. It is clear that in an ITS, such features of the knowledge domain play a crucial role in determining tutorial style (see for example, our earlier mentioning of the WHY system). Since knowledge of musical structures is necessarily uncertain and ambiguous, tutorial styles which rely on simple imparting of beliefs to the student are necessarily precluded. One adequate response to developing a dialogue model which is appropriate to tutoring in such uncertain knowledge domains, is the use of dialogue structures which we term *critical argument*. The model will be described in detail in the next section, but for now it suffices to state that a critical argument is precisely the kind of discussion which one would expect between an academic and student concerning the contentions of the latter, in which the former is not seeking to impart knowledge but rather to refine beliefs and produce a critical attitude on the part of the student towards their own and others' beliefs.

As was discussed earlier, to say that uncertain knowledge domains may require a specific kind of dialogue structure, is not to claim that these dialogue structures are restricted to such domains. In our dialogue model, different tutorial strategies can be characterised by consistent choice of specific kinds of *dialogue goals* (eg, questioning, explaining, and so on),

and it is these specific *tutorial strategies* within the dialogue model which may be more or less appropriate for various knowledge domains. In other words, part of the generalisability of a dialogue model resides in its supporting a variety of teaching strategies.

An example domain: knowledge of musical structures

The model of knowledge of musical structure embodied in the ITS was influenced by issues of cognitive validity and of how best to represent the knowledge which a human expert in analysis of musical structure possesses, in a manner which is most suitable for explicit communication to a student (Wenger 1987). A simple approach to representing knowledge of musical structures would be for a human analyst to prestore this information for a set of pieces in the tutoring system. We have not adopted this approach for a number of reasons.

- as a pragmatic point, human analysts themselves often disagree over analyses,
- at present there are few formalisable and explicitly agreed principles in musical analysis;
- we want the system to be able to deal with melodies input and chosen by the student (within limits of the 'musical language' which the system can handle), without having to rely on prestored structural analyses.
- prestored facts alone lack explanatory power: we would like the system to derive structural (or 'grouping') analyses, so that we can use the history of the computation for explaining the occurrence of musical structures.

The method used to derive grouping structures for tonal melodies uses frame-based schemata for stereotypical grouping structures. The grouping schemata are used to guide a chart parser (described in Baker 1988a and Baker 1989, in press), using a grammar of chord function to identify 'well-formed' progressions which indicate phrase boundaries. It is presently restricted to grouping structures at the phrase-level for very restricted musical genres, but some indication as to further work on sub-phrase levels has been given. From the point of view of generating tutorial dialogues, it is important that the system will not necessarily derive a single correct answer as to phrase structure: a set of hypothesised grouping structures

will be derived, together with their attendant justifying factors in terms of harmony, recurring parallelisms, or patterns (Lerdahl & Jackendoff 1984) and related features on the musical surface. These surface features correspond to the Lerdahl & Jackendoff 'grouping preference rules', in terms of Gestalt discontinuities in rhythm, dynamics, and so on. In addition, the system has a set of global constraints on phrase boundary sets, derived from Lerdahl and Jackendoff's well-formedness rules. Two very simple 'constraints' on combinations of beliefs which the system may possess are

(1) approximately evenly spaced grouping boundaries are preferred;

(2) very small distances between boundaries are dispreferred.

The knowledge of musical structure available to the ITS can therefore be summarised as in figure 2.

Figure 2 Knowledge of musical structure available to the system

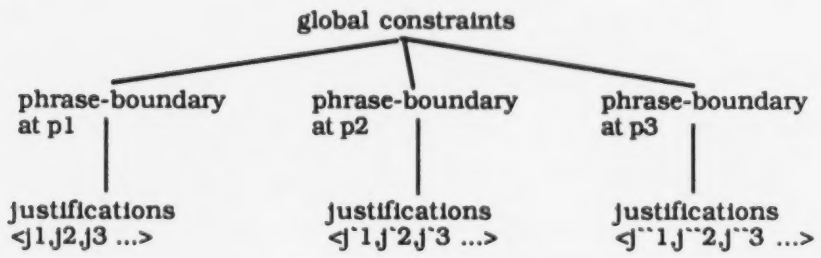

An example would be

belief 1 - 'phrase boundary at point p1'

j1 - 'exact match in harmony (vi IV ii) in note sets n1 and n2'

j2 - 'exact match in pitch contour, reductional level 3'

belief 2 - 'phrase boundary at p2'

j`1 - 'interval leap of major 6th'

j`2 - 'match at reductional level 2, upward transposition of major 3rd'

The basic tutorial scenario is that the system or student makes some *claim* concerning the

presence of musical structures called 'phrase boundaries' in the melody being considered (eg "there is a phrase boundary at point p1"), and a critical argument dialogue ensues concerning the reasonability of this claim. Suppose that the student asserted a boundary close to belief 2, which divided the melody into uneven sized groups. A global constraint 'prefer boundaries of even size' could be invoked, and the student would be invited to critically discuss their reasons for claiming that there is a phrase boundary at this point, in opposition to the system's belief set. To simply present the set of beliefs concerning phrase boundaries and their justifications - perhaps in the form of simple text formats - would not be an effective basis for effective tutorial strategies, since no account is taken of the student's existing knowledge, learning goals and abilities, and ability to integrate this knowledge with existing knowledge. As will be argued in the next section, the system's knowledge must be communicated in the form of an interactive tutorial dialogue, which achieves tutor and learner goals, is focussed and is relevant.

A model for generating human-computer critical arguments

In the previous section we discussed the nature of dialogue models in Intelligent Tutoring Systems, the form which a dialogue model should take, and the close relationship between domain and dialogue models. We also introduced an example cognitive model for the knowledge domain of musical structure. In this section we shall describe a dialogue model which is based on this example knowledge domain. The discussion will be largely descriptive, critical discussion of the model and its relationship to existing work in human-computer dialogue research being dealt with in the final two sections of the chapter. The structure of this section is as follows:

- remarks on methodology
- higher level knowledge representations for dialogue
- a dialogue model for critical argument
    - a typology and discussion of dialogue goals represented in the model
    - components of dialogue goals (goals, preconditions, effects and actions) and their functional role

- an example dialogue and its derivation from the model
- the algorithm for dialogue generation, discussed with reference to the example.

Remarks on methodology

First, we need to make a few remarks about methodology. We have used approaches common to cognitive modelling techniques; which are to consider available empirical and theoretical research on the structure of dialogue, and to attempt to formalise and synthesise this existing research in the form of a formal computational model. A further methodological technique is the use of *protocols* of tutorial dialogues between human teachers and students (see, for example, Collins' 1977 work on the WHY system). The number of available theoretical models of dialogue structure which we could use to *interpret* such data is, however, very small. As Petrie-Brown states (elsewhere in this volume), at the time of writing (1988) the work of Power (1979) is the only extant computational model of dialogue. We would therefore maintain that with our present paucity of theoretical knowledge, it is impossible to produce theory 'bottom up' from the data, but that protocol analysis may be used as a weak test of the plausibility of a model (although it would not necessarily distinguish between competing theories). The research described here therefore forms the first stage in an iterative research program, of hypothesising computational models of human-computer dialogue, and testing the successive prototype models with protocol analysis and empirical evaluation of the effectiveness of its teaching ability.

Representing the domain within the dialogue model

The knowlege representation discussed in the previous section represented the listener's ability to consciously articulate the result of the listening process, in the form of a set of *beliefs* concerning musical structures, and to further intuit a set of *justifications* for those beliefs. However, no set of beliefs exists in a 'conceptual vacuum', without integration with an individual's existing knowledge in that domain. This implies that we need some higher level representation for these conceptual connections between beliefs and their possible justifications. When these beliefs come to be expressed in coherent dialogue, we additionally

need some mechanism for characterising dialogue topic or focus, in terms of this higher level representation of conceptual dependency. In other words, a simple representation of belief and justification sets is not a sufficient basis for generating dialogues, since we have no representation of the concepts which interrelate possible justifications for belief, no conception of topic and focus, and hence no means of relating topic to tutorial goals. The higher level knowledge representation which we have chosen is that of *semantic networks* (Quillan 1968). This representation is particularly suitable here, since it has been used in existing research for models of memory recall (Anderson 1984) and for characterising dialogue focus (Grosz 1977).

In addition to possessing knowledge in a domain, a dialogue participant needs to know how to *use* that knowledge in interactive dialogue. So for a student to be able to take part in a critical discussion in some knowledge domain, at least two kinds of knowledge are needed:

(1) knowledge of factors which are relevant to supporting claims, and

(2) knowledge of how to advance claims and their support in the form of a coherent argument, perhaps in opposition to those put forward by another dialogue participant.

In order to learn the factors which are relevant to supporting claims concerning musical structures a student will need to understand the *concepts* which interrelate different justifying factors, and to be able to integrate these concepts with those which they have already learnt.

Now clearly, given that the system has sufficient conceptual knowledge in this domain, we must assume a certain theory of how these concepts could be *learned* by the student as a result of engaging in the dialogue. Of necessity in this knowlege domain, methods of teaching a concept ('phrase boundary ', for example) by presenting a set of positive, negative and 'near miss' examples (Winston 1977) are precluded - the system has no examples of which it is certain, and must argue for each case. One possibility for tutoring such concepts is to tutor the concepts which underlie the justifications for positing instances of these concepts - i.e. tutoring a concept by tutoring the concepts to which it is most closely related. For example, suppose that three justifications for an instance of a phrase boundary are

(i) presence of a harmonic cadence,

(ii) a change in dynamics, and

(iii) a harmonic parallelism.

Clearly, the concept of what a cadence is depends on understanding of many others, such as 'chord', 'dominant chord' and others which may need to be tutored as precursory knowledge. For example, with second justification, what is important is that it is just one instance of contrast, and that if the student can be taught this, their knowledge can become sufficiently generalised so that they will be able to understand that any *other* form of musical contrast would similarly be a justification here. Figure 2 shows a portion of a semantic network (derived from existing musical theory) attached to each type of assertion which the system can make, and which interconnects its set of possible justification types.

Figure 3

Part of a semantic network relating justifications for possible instances of phrase boundaries

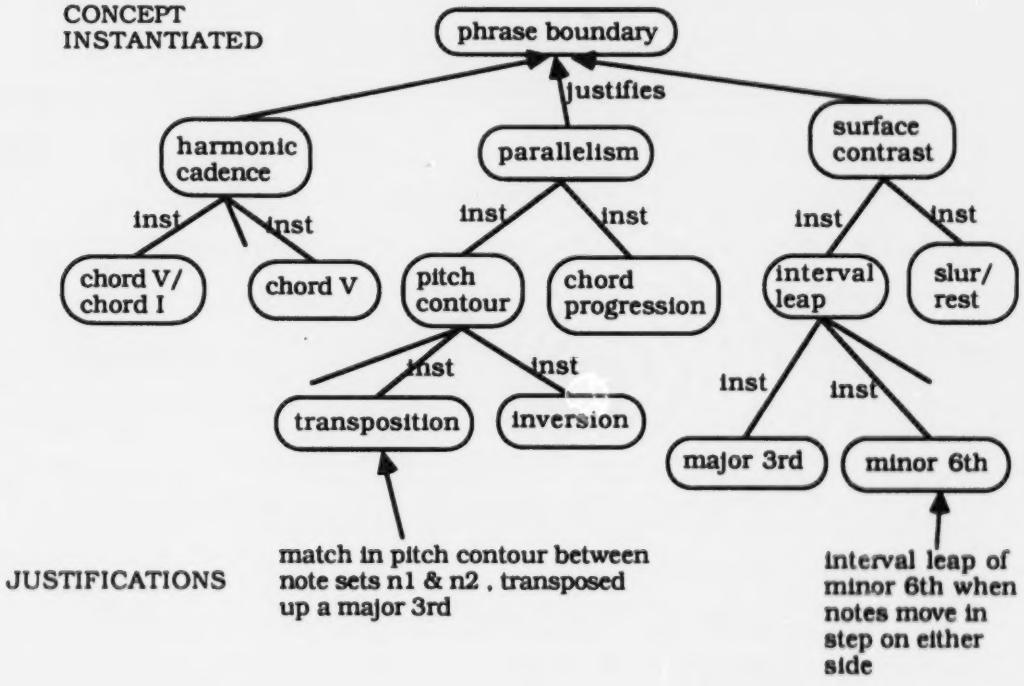

In an ITS, the knowledge representation chosen is usually crucial in influencing a number of features of the system. In this case the semantic network will have the following roles in the model:

(i) *choice of tutorial goal* - since the network connects closely related concepts,tutorial goals will be chosen which are connected to previous goals (the goals of teaching that particular concept);

(ii) *user modelling* - the student's knowledge can be modelled as a simple subset of the tutor's network (we are not addressing the problem of the student possessing or acquiring a quite different conceptualisation of the domain);

(iii) *dialogue focus* - this relates to choice of tutorial goal, since current dialogue focus will relate to the present position in the network, and dialogue focus shifts will be made smoothly and relevantly to concepts which are implicitly in focus through their connection to the current focus in the network.

In the network shown above, for example, if we are currently teaching the concept 'interval leap', then the concepts 'interval' and 'musical surface' are in implicit focus. It would be possible to make the new tutorial goal 'teach interval', with its attendant focus shift. As it stands, such a network could be used in a straightforward fashion (given the reasonably large knowledge engineering effort required to represent it in the first place) to generate explanations for system beliefs by exploiting links in the network, omitting those which we think the student already knows (Weiner 1980) and those which we have not tutored yet, which the student can not integrate with their existing knowledge. However, this would give little coherent dialogue structure, since the explanation would consist simply of pairs of question/explanation responses. Current research in dialogue and explanation (Draper 1987; O'Malley 1987) has moved towards a view of explanation as an emergent property of an ongoing cooperative dialogue, rather than as a kind of discourse genre, where an explanation can be given in the form of a single piece of generated text. Explanation - and, we would argue, *learning* - is something which emerges from a dialogue.

The dialogue model

The dialogue model is based on viewing the kinds of 'conversational move' (Reichman 1985) which might be made in a critical argument as possible *dialogue goals* to be pursued by each participant. Dialogue generation is therefore viewed as generation of a *plan* of utterances which aim to achieve desired effects on the 'world state' of the beliefs of the other participant. For the ITS, therefore, tutoring is fundamentally a matter of planning to achieve certain effects on its representation of the student's beliefs - the student model. Dialogue goals are represented as STRIPS (Fikes & Nilsson 1971) operators, which represent possible subgoals, preconditions for goal pursuit, assumed effects on belief sets, and *actions* (dialogue utterances). A special feature of the dialogue planner is its combination of *top-down* and *bottom-up* elements in structuring the interaction. Top-down elements include decisions as to which concept should be taught, and with what type of dialogue goal it should be pursued. Bottom-up elements include the suggestion of subtopics to be taught at any point in the dialogue (DISCUSS(related c) in the goal tree of figure 4), and focus-based constraints on topic shifts. Constraints on possible relevant focus shifts are determined by a focus model based on *spreading activation* (Anderson 1984) from the current focus in the semantic network. We shall discuss theoretical models of focus in the next section. For the present it is sufficient to understand that a record of the current topic under discussion and its position in the network is abstracted from the dialogue history. New topics which are mentioned are passed to a program which propagates these appropriate new activation levels. Topics which are possibly relevant are those which are sufficiently *active*, in terms of spreading activation to other concepts in the semantic network. A distinction is made between *implicit* and *explicit* dialogue foci: explicit topics are all those which have been explicitly mentioned (eg 'phrase boundary', harmony, interval), and implicit foci are those which have recieved sufficient activation in virtue of their connection to explicit foci in the semantic network. Within explicit foci, a single focus (the topic determined by the planner) is selected as the *currently active focus*. Global dialogue coherence will thus be achieved in virtue of sets of utterances being part of the same coherent plan, whilst more local coherence is achieved in terms of focus shifts in the network.

Let us examine the model in more detail in terms of specific representations of goals, algorithms for dialogue generation, and discussion of a hypothetical dialogue which could be generated.

Dialogue goals for critical argument

Possible dialogue goals, for system and student to pursue, are represented in terms of a hierarchy of abstraction. Only the terminal nodes of the goal tree have actual dialogue actions attached. Figure 4 shows such a goal tree, for discussing some concept 'c'.

Figure 4
The hierarchical tree of dialogue goals

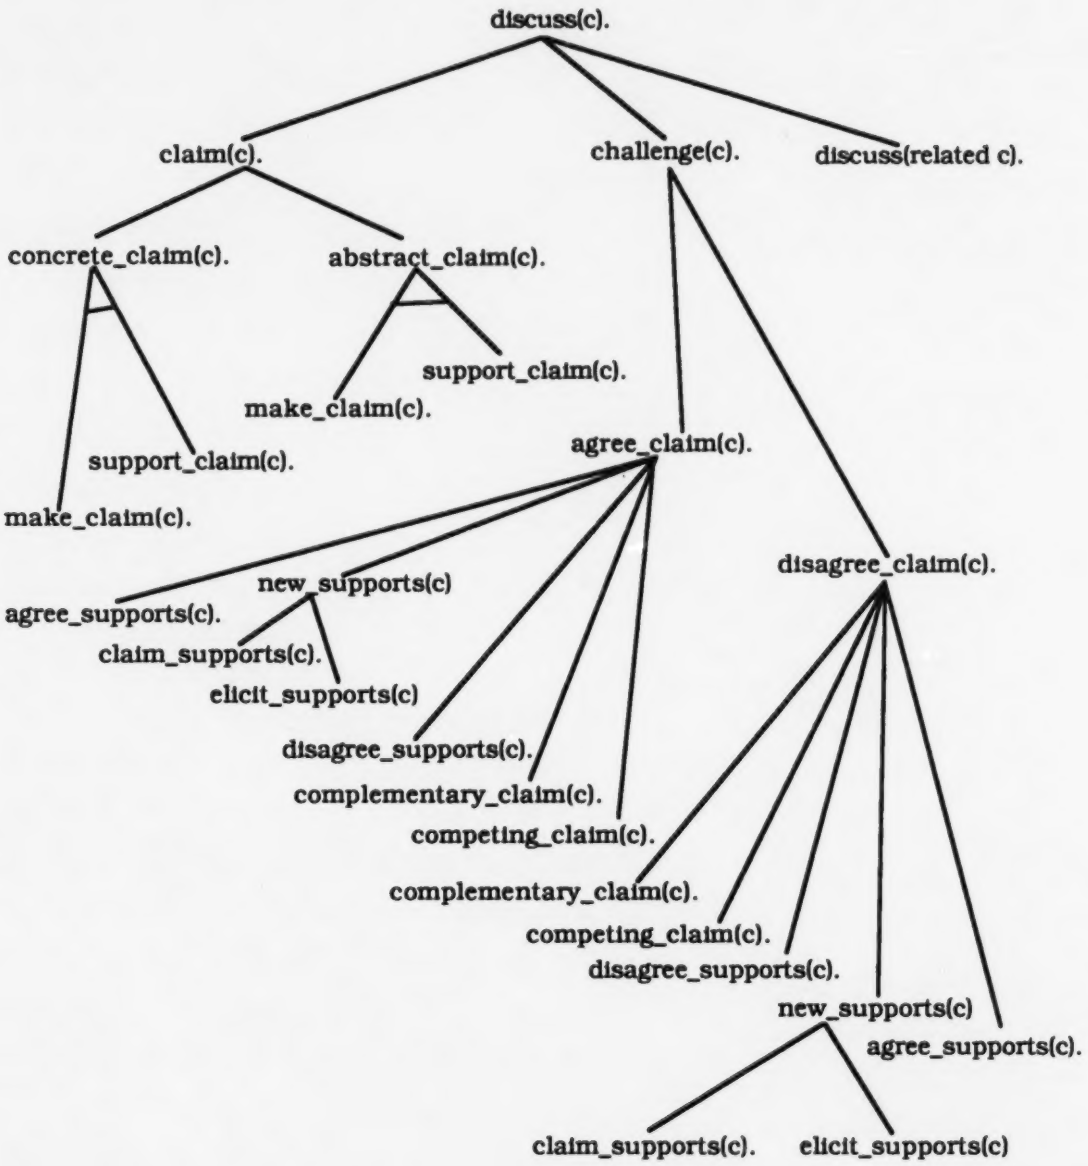

Higher level planning decisions

The most general dialogue goal is to DISCUSS a particular concept. The decision as to which concept to discuss would be decided by a higher level *topic planner*, which would plan a sequence of concept teaching goals in accordance with the student model. Such a planner has been described by Peachey & McCalla (1986), which depends on representing the subconcepts upon which a concept depends. The top level planner embodied in our system can use the conceptual relations embodied in the semantic network to achieve a similar effect (Figure 3). In addition, suggestions as to which concept to tutor can be generated 'bottom-up', from concepts which are closely related to the one being currently taught (according to the top-down concept plan). In human terms, this corresponds to a teacher having an initial fairly general plan of topics to be taught, but being willing to modify that plan, should the opportunity arise to tutor relevant subtopics, which may be suggested by the student, for example. There are a number of interesting research problems involved in developing such a topic planner, which we shall not elaborate here (since our primary concern is with dialogue planning). These include issues as to when a plan should be abandoned (and replanning take place), or whether a 'local fixup' should be done to retain the original plan. In addition, there is the general planning issue of attempting to satisfy multiple simultaneous goals (of the teacher and learner).

We stated in introduction that *tutorial styles* would correspond to consistent choices of specific kinds of dialogue goals. This means that, in addition to a higher level topic planner, we need a further decision mechanism to choose an appropriate sequence of *tutorial styles* with which to teach each concept, which takes account of the student's preferred and most effective learning styles. As will emerge, choice of tutorial style does not only involve choice of dialogue goal, but also of *speaker roles* to be adopted in pursuit of that goal. For example, consistently choosing the CLAIM goal with the *teacher* as speaker would correspond to something like simple 'instructing'. With the teacher adopting the goal of the *student* making the claim, however, this is closer to 'elicitation'. At present no tutoring systems have addressed the problem of choosing an appropriate rational sequence of different tutorial styles, a single exception being the DOMINIE system (Spensley & Elsom-Cook 1988). In DOMINIE, teaching

styles are determined using a set of rules, whose antecedents refer to the student and dialogue model states, and whose consequents refer to teaching styles. We shall use a similar set of rules, where each possible teaching style is characterised as a specific combination of dialogue goals and speaker roles - i.e a labelled path through the dialogue goal tree. Clearly, a problem of conflict resolution may arise, and we want to resolve these conflicts in a non-arbitrary way, which is educationally principled. As Elsom-Cook has suggested (1984) one way of characterising alternative tutorial styles is in terms of their *degree of directedness*.The Guided-Discovery approach to Intelligent Tutoring can be viewed in terms of the strategy of "decreasing intervention" (Elsom-Cook 1984) , where the system uses the least directive strategy which is possible (i.e. whose conditions are satisfied) at any point in the interaction, with an overall goal of steadily withdrawing support towards a position of nil directiveness. Since our dialogue model is designed to support guided discovery, we propose to implement this by associating a (numerical) degree of directedness with each tutorial strategy (rule rhs). A principled conflict-resolution strategy can therefore be based on degree of directedness, as determined by the desired overall educational approach.

Making claims

A fundamental distinction is made between dialogue goals which involve the making of some type of CLAIM, and those which involve CHALLENGING a claim which has already been made. In addition, there is the possible subgoal of discussing a related concept (a 'bottom-up' suggestion of a teaching topic). Note that a claim which is challenged by an initiator ('I') need not have been made by a different conversant, the respondent ('R'): it is conceivable that for didactic purposes, 'I' may wish to make a putative claim and subsequently challenge (or 'examine') that same claim, supposing that 'R' lacks sufficient knowledge to make a claim and critique it. So doing would correspond to a possible tutorial style which might be termed 'annotated example'.

At the next level of abstraction, CLAIM goals divide into either *abstract* or *concrete* claims. An abstract claim corresponds to something like an 'explanation', and involves making some

statement about what would justify a hypothetical instance of a concept, without actually stating some concrete instance. Concrete claims involve making an 'illustration' of a concept by describing a specific instance and its justification. With respect to discussion of a specific topic, this is related to Reichman's (1985) distinction between 'instance' and 'issue' context spaces. For either a concrete or abstract claim, the subgoals are simply

- the making of a claim (MAKE_CLAIM(concept)) and
- giving a believed set of justifications for that claim (SUPPORT_CLAIM(concept)).

In the goal tree, these subgoals are marked as *conjunctive*, which means that the conversant (system or user) is committed to producing supports for any claim. This is one way of avoiding the necessity to *elicit* supports on this branch of the tree (see the subsequent discussion of components of the challenge goal). All other alternative goals are disjunctive in a *non*-exclusive sense (i.e. any combination of the subgoals may be chosen, in a specific ordering from left to right, for tutorial reasons).

Making challenges

On the CHALLENGE branch of the goal tree, the goals divide into whether the current conversant agrees with, or disagrees with, the claim which is to be challenged. In the case of a participant challenging a claim which they themselves have previously made, with which they presently disagree, (according to the dialogue history), this corresponds to simply having revised their beliefs (we shall describe the mechanism for belief revision subsequently). Where the challenging participant agrees with a claim made by another, but decides to challenge, this corresponds to a decision to adopt a tutorial style of 'critiquing' the grounds for making such a claim. An interesting property of the subgoals for challenging a claim, with either agreement or disagreement, is that the subgoals are exactly the same in each case: they differ only in their *ordering*. There are intuitively acceptable reasons for this. In critical argument conducted with a view to didactic purposes, one would wish to examine a student's reasons for making some claim whether one agreed with the claim or not. Secondly, in human critical argument where one adopts an adversary position, it seems most natural to begin by discussing those points on which one disagrees initially, and move on to points of agreement afterwards (or else assume that what is unsaid is agreed) The converse seems to be true in situations where one

shares a conversant's beliefs. Weiner (1980) has discussed these issues with respect to generation of explanations from expert systems. Many of the subgoals here are straightforward:

- a conversant states those justifications with which he agrees (AGREE_SUPPORTS),
- with which he disagrees (DISAGREE_SUPPORTS) and

in cases where a conversant believes that there is some possible justification which the opponent has not stated (NEW_SUPPORTS (c)), there is a choice of either

- stating those justifications which he believes which the other conversant has not stated (CLAIM_SUPPORTS), or
- attempting to elicit unstated supports for a claim (ELICIT_SUPPORTS(c)).

In addition, there are two further kinds of subgoal which bear close resemblance to the 'rules for Socratic tutoring' described by Collins (1977). In his work he discussed such tutorial strategies as 'probing for necessary factors' and 'setting up counter-examples for insufficient factors'. Whereas they were discussing factors influencing rainfall, we are concerned with justifications for belief, but many of the required dialogue moves are similar. In cases where the student has stated that a set of justifications justify an instance of a concept at some position then the system would look for other positions in the melody (perhaps already encapsulated in its own belief set) where the same justification set also holds, and ask the student if they would also consider this to be an instance of the concept (COMPLEMENTARY_CLAIM), possibly leading to a revision of their beliefs. In a knowledge domain where certain beliefs could be shown to logically or practically contradict each other, an adversary in a critical argument would clearly want to point this out . In our domain there are no such strictly logical contradictions, but rather only 'contradictions' between certain sets of beliefs in terms of a number of separate principles - the 'global constraints' on phrase boundary sets, which were mentioned earlier (the section 'An example domain: knowledge of musical structures'). For tutorial purposes this means that if the student makes a claim concerning a phrase boundary which would produce unevensized or very small groups in combination with a previous claim of their own, or the system's, or in terms of a system belief which has not yet been stated, then the subgoal of stating that COMPETING_CLAIM could be

pursued.

Components of dialogue goal operators

(i) actions

The terminal nodes of each branch have actual dialogue actions attached. With considerable conceptual economy, we can view *all* dialogue actions (other than those presequence adjacency pairs attached to preconditions, which are explained subsequently) in terms of a *single* action- *INFORM(concept)* - whose argument 'concept' can correspond to

- the participants beliefs concerning an instance of a concept,
- its possible justification, or beliefs concerning the agreement or disagreement with some other belief.

As will be discussed in the next main section, INFORM can be viewed as an example of a *dialogue game* (Levin & Moore 1977), and accordingly has

- an attached text generation template (the action)
- representation of current initiator and respondent (speaker roles), and
- effects (represented in the main body of the goal operator).

The provision of alternative speaker role variables attached to the INFORM dialogue game effectively preserves the dialogue *symmetry* which was discussed in the introduction: both the system and the student can make INFORM statements, with correspondingly symmetrical effects.

(ii) Effects

The primary effects of dialogue actions relate to speaker roles. One simple effect of performing the INFORM action is that the belief uttered by the initiator is assumed to be *remembered* by the respondent. We use the term 'remembered' at this point since we are presently not addressing the complex issue concerning whether an imparted belief takes the form

believes(R believes (I b1))

(respondent believes that initiator believes that b1)

or simply

believes(R b1)

(respondent actually believes that b1).

When the student is initiator, this involves updating the system's belief set (given certain conditions for acceptance, discussed later), and in the case of the system initiating, this involves updating its student model. Other kinds of effects of an action are those essentially involved in 'bookkeeping': updating focus records (increasing activation of certain mentioned topics) and updating the dialogue as well as the student model.

(iii) preconditions

Most of the theoretical burden in the goal representation falls upon attempts to satisfy *preconditions*. Preconditions are those factors of the current dialogue state which need to be checked as *satisfied* before the goal is pursued and dialogue actions eventually performed. Such features of the dialogue state relate to

- *agreement* on mutual cooperation in pursuit of the joint dialogue goal, its topic, and speaker roles
- the knowledge possessed by dialogue participants (a participant can't pursue a topic concerning which they have no beliefs)
- active dialogue foci, and relevant focus shifts

For example, the STRIPS operator for the CLAIM goal is as follows:

goal: CLAIM(concept)
subgoals: abstract_claim(concept) concrete_claim(concept)
preconditions:
p1: PREANNOUNCE(Initiator Respondent).
p2: PREANNOUNCE(topic = concept).
p3: PREANNOUNCE(goal).
p4: active-focus(concept).
p5: tutorial-goal(concept student-model).
effects:

[set new roles, focus, dialogue history]

action: nil

Suppose the tutorial strategy decision and tutorial goal planner has decided to make a CLAIM concerning this concept. At this level, preconditions simply check whether this topic, dialogue goal and set of speaker roles are accepted by the respondent. In addition, it checks that discussing this concept at this point represents an acceptable focus shift, and that such a claim represents a reasonable learning goal for the student according to the student model. We shall discuss how such preconditions are *satisfied* subsequently. Assuming that preconditions can be satisfied for this goal, further assume that the following abstract claim goal is chosen:

goal: ABSTRACT_CLAIM(concept)

subgoals: make_claim(concept) support_claim(concept)

preconditions:

p1: PREANNOUNCE(INFORM(instance concept) student-model).

p2: exists(instance(concept) belief_set).

effects: nil

action: nil

The first precondition is designed to avoid telling the respondent what (s)he already knows: it is automatically marked as satisfied unless, according to the student model, the dialogue initiator is unsure whether the respondent already knows the information which is about to be imparted. The second precondition simply checks that there exists an instance of the concept as part of the initiator's belief set which can be imparted. Neither of these higher level goals have associated actions: action components are contained within the subgoals - for example the action for make_claim(concept) is INFORM(Initiator Respondent instance(concept position-1)). The effects of the action on the student are assumed to be simply that the belief which has been communicated is remembered; but for more abstract goals there are of course effects on focus and the system's overall model of the interaction - student model, dialogue

history (an abstract history of claim types, speaker roles and topics) and dialogue model (current speakers, focus and so on).

Preconditions therefore play a crucial role. They are of two kinds: those which require dialogue adjacency pairs to be generated in their attempted satisfaction, and those which do not. In implementational terms, checking preconditions first which lacked attached dialogue procedures could reduce computational effort, supposing that these preconditons were to fail. We have mentioned several non-dialogue preconditions, relating to

- checking suitability of learning goals
- checking acceptable focus shifts
- checking that requisite beliefs exist to be communciated (for example, that a concrete instance exists to be stated, that a possibly complementary claim exists, that unstated justifications exist, and so on)
- checking that a suitable previous dialogue goal type exists in the dialogue history (for example, that a previous claim exists to be challenged)

Preconditions which have associated pieces of subdialogue have the purpose of sustaining a phase of *negotiation* between speakers concerning mutually acceptable dialogue goals to pursue. Here we have adopted the somewhat unnatural approach of *explicit* negotiation in the form of *preannouncements* (as adopted by Power 1979), concerning

- what topic should be discussed,
- what type of dialogue goal should be pursued,
- who should adopt roles of intiator and respondent in pursuit of that goal,
- avoiding stating what is alreadty known

For each such precondition, there is an associated *adjacency pair* (Schegloff & Sacks 1973) which specifies the form of the initiator's utterance and possible responses. Each adjacency pair is generated by the PREANNOUNCE procedure, which has an associated text format, arguments which relate to the specific precondition, speaker roles, and a set of associated possible responses. The PREANNOUNCE format can be a piece of text as trivial as

"Shall we _ "

whose arguments can be such things as

"talk about <concept>"

"make a <goal>"

"begin by <initiator> speaking first".

Very simply, there are only two allowable responses - "OK" and "NOT OK". In the first case the precondition is satisfied, and in the second not. If any precondition is not satisfied, its parent goal *fails*, otherwise the goal succeeds, its subgoals are pursued, and any attached actions are incorporated into the plan to be conversed ('interpreted').

We stated earlier that these preconditions have the function of attempting to secure mutual cooperation and acceptance for the dialogue goal to be pursued, by means of explicit negotiation. Such preconditions are not only pursued by the ITS, but also by the student in an attempt to secure the *system's* cooperation. This results from the fact that the dialogue goals made available to the system, are also available in their entirety to the student: in other words, the possible dialogue actions available are *symmetrical* between system and student (we shall discuss the implications of this concept in conclusion). In the algorithm for generating dialogue from such a goal tree, the system makes its precondition utterances using text formats. When the student is attempting to secure agreement for a dialogue goal (eg the student wishes to CHALLENGE an utterance made by the system), the preconditions for goal pursuit are presented to the student in the form of menu choices to fill in valuse in the same (symmetrical) formats, which the system can similarly respond to with acceptance or rejection. For example, in pursuing the CHALLENGE goal, the student may attempt to satisfy the topic PREANNOUNCEMENT precondition by choosing the topic 'harmonic cadence' from a range of presented alternatives. It may be that the system would choose to reject this topic as an unsuitable learning goal and in terms of focus shifts, in which case it could respond

"I do not wish to pursue this topic now because

1. I do not think you would be able to understand it at this point in our discussion

((checking student model))

2. it is not directly relevant to what we have just been discussing"

((checking concept activation in focus model))

and the student would be invited to choose a new dialogue goal, which the system may or may not subsequently accept.

We propose to conclude this discussion of the dialogue model by considering a hypothetical fragment of dialogue and describing algorithms which could generate it.

An example hand-simulated dialogue

<system = s; user = u>

Hypothetical dialogue	Commentary
	((system goal choice))
	((topic goal = phrase boundary))
	((tutorial style = INSTRUCT))
1 s: I want to propose what we discuss next.	**((focus & learning goal preconds))**
2 (OK? QUIT)	
3 u: OK	
4 s: Why don't we talk about phrase boundaries?	**((topic precond))**
5 (OK? or NOT OK?)	
6 u: OK	
7 s: ... I want to begin speaking ...	**((speaker roles precond))**
8 (OK? or NOT OK?)	
9 u: OK	
10 s:..and to make a claim...	**((goal precond = CLAIM))**
11 (OK? or NOT OK?)	
12 u: OK	
13 s:... about a specific instance of a phrase boundary...	**((goal precond=CONCRETE_CLAIM))**
14 (OK? or NOT OK?)	
15 u: OK	
16 s: ... in the last ten notes of the melody ...	**((preannouncement precond))**
17 (OK? or NOT OK?)	
18 u: OK	
19 s: I believe that there is a phrase boundary	
20 at position p6....	**((action MAKE_CLAIM))**
	((p6 indicated on graphic interface))
21 (CONTINUE? INTERRUPT?)	
22 u: CONTINUE	
23 s:...because	
24 there is an interval leap of a minor 6th	**((action SUPPORT_CLAIM))**
	((topic INTERVAL activated))
25 (CONTINUE? INTERRUPT?)	
26 u: CONTINUE	

27	s: and, a harmonic change from chord V to I.	((action SUPPORT_CLAIM)) ((topicsHARMONY, CHORD activated))
28	Would you like to propose what to discuss next?	((user goal choice))
29	(YES? NO?)	
30	u:YES	((user accepts goal choice possibility))
31	s: Available conversational moves are	
32	CHALLENGE	
33	CLAIM	
34	ELICIT	(('canned' explanations available. System has option of suggesting the next goal for the user to choose))
35	u: CHALLENGE	
36	s: OK, that's fine by me..	((dialogue goal accords with system choice of tutorial style))
37	s: Available topics are	
38	PHRASE_BOUNDARY HARMONY CHORD	
39	INTERVAL MINOR_CHORD	((user topic goal choice & attempted satisfaction of topic precondition))
40	u: PHRASE_BOUNDARY	
41	s: OK	((system can satisfy focus shift precond))
42	Who would you like to INITIATE the CHALLENGE	
43	about PHRASE_BOUNDARY?	
44	(ME? SYSTEM?)	((speaker roles precond)) <note that the user could have chosen that the SYSTEM should CHALLENGE its own previous CLAIM>
45	u: ME	
46	s: OK	
47	Which claim would you like to CHALLENGE?	
48	(select the claim from the dynamic claim menu)	
49	u: PREVIOUS	
50	s: Available conversational moves to choose are	
51	AGREE_CLAIM	
52	DISAGREE_CLAIM	((canned text explanations given))
53	u: DISAGREE_CLAIM	
54	s: OK	
55	Available conversational moves to choose are	
56	COMPLEMENTARY_CLAIM	
57	COMPETING_CLAIM	
58	DISAGREE_SUPPORTS	
59	CLAIM_SUPPORTS	
60	AGREE_SUPPORTS	((canned explanations available))

61 u: COMPLEMENTARY_CLAIM
62 s: OK
63 Indicate on the interface the position where you believe that the justifications
64 there is an interval leap ...
65 and, a harmonic change from chord V to I
66 are also true.
67 u: <marks position p9>
68 s: thankyou for telling me that:
69 however, to have a phrase boundary here contradicts a general principle which I believe,
70 which is that phrase boundaries are evenly spaced. Since I believe that there is a phrase
71 boundary also at p4, I do not want to add this to my set of beliefs.
72 Choose another goal?
73 (YES? NO?)
74 u: YES
75 s: Available conversational moves to choose are
76 COMPETING_CLAIM
77 DISAGREE_SUPPORTS
78 CLAIM_SUPPORTS
79 AGREE_SUPPORTS
80 u: CLAIM_SUPPORTS
81 s: Possible supports here are
82 INTERVAL_LEAP CADENCE PATTERN_RECURRENCE(etc.)
83 Which of these do you claim is also true at the position I indicated?
84 u: PATTERN_RECURRENCE
85 s: what type of pattern?
86 RHYTHM? MELODIC_CONTOUR?
87 u: rhythm
88 s: indicate the two sections on the interface.
89 u: <indicates>
90 s: Thankyou: I agree that this is also a possible justification here and I will add it to my set
91 of justifications for this belief about phrase_boudaries.
92 Choose another goal?
93 (YES? NO?)
94 u: NO
95 s: I want to propose what we discuss next. **((system goal choice))**
96 (OK? QUIT)
97 u: OK
<<continues>>

The algorithm for dialogue generation

In the light of our previous discussion, the commentary is reasonably self explanatory, and the simple underlying algorithm for dialogue generation can now be easily understood. At the most general level, the dialogue structure is generated by iterations through a cycle of

- *choose a goal* to pursue (topic & dialogue goals)
- *generate a plan* to achieve the goal
- step through the plan and *converse* according to it.

The iteration terminates when either all higher level (topic) goals are marked as achieved or when the student terminates the dialogue voluntarily (types QUIT). For the dialogue planner which we have described, goals are chosen according to a higher level plan of topics and tutorial styles. For example, in the hypothetical dialogue above, the first goal chosen was the topic 'phrase_boundary' with a tutorial style of instructing (imparting some belief with a CLAIM dialogue goal). Recall that choice of dialogue goal is further constrained by relevant focus shifts and our model of the student's preferred learning style. The dialogue plan is simply a list of dialogue actions paired with those preconditions which require presequence adjacency pairs to satisfy them, rather than simply a plan of *actions*. This reflects the fact that the plan should include all and only those *utterances* which need to be generated, and in this case this includes a number of presequences. In the dialogue above, the plan for the first section - where the system chose a CLAIM dialogue goal with a PHRASE_BOUNDARY topic - would be

```
(  ((PREANNOUNCE(topic phrase_boundary)
     PREANNOUNCE(initiator  system)
      PREANNOUNCE(goal  CLAIM)
     PREANNOUNCE(instance phrase_boundary  position))
                                                    nil)
   (nil
   (INFORM(instance phrase_boundary position)
   (INFORM((supports (interval_leap  value)
                        (harmonic_cadence V_to_I)))))
```

In generating the plan, preconditions which do not require dialogue to be generated are first checked. If they fail, then the plan fails and the higher level planner must choose new goals. Where a goal has disjunctive subgoals, all these subgoals are incorporated into the plan, with their attendant preconditions. For the present we have adopted the convention that generation of a plan for conversing *alternates* strictly between system and student (note that this

alternation of goal choice is quite independent of speaker *turn taking* since on successive iterations a conversational initiator may choose for the other conversant to initiate the dialogue - eg choosing that the other should make a claim, or posessing the goal of 'eliciting' a claim). We have described the procedure for the system's generation of its plan, but the system creates an identical type of plan when the student is choosing the goal. It does this by simple top-down presentation of available dialogue and topic goals, stepping through the goal hierarchy. Thus lines 28 to 94 in the dialogue above correspond to the system allowing the user to create a dialogue plan. The only fundamental difference lies in the way in which the plan is *conversed*: when the system is conversing a plan, it does so according to text generated from templates; but the student's means of conversing with the system is mediated via choice amongst available arguments for the same text templates attached to preconditions which are presented to the student (for example, in lines 40 to 45 in the example dialogue).

So far we have described how a plan is generated from a goal where the plan is that of the system or of the student. Conversing the plan fundamentally involves stepping through the plan and executing it - i.e. making the appropriate utterances. If the text corresponding to a precondition presequence receives the response 'NOT OK', then its goal fails, and the action component is not executed (by system or student) . If it receives the response 'OK' then the goal succeeds, the action is executed, and the *effects* of the action are propagated. Where the action involves imparting some beliefs, certain straightforward effects (recorded in the student and dialogue models) are that the belief is remembered by the student (or the system, if it accepts the belief), and focus shifts are recorded. In the execution of disjunctive subgoals (simply, where there is a list of actions for a plan element), then after each action is sucessfully executed, the option of *interrupting* the conversation is provided for the respondent. For example, suppose the system were to interrupt after the student had uttered the COMPLEMENTARY_CLAIM goal action in line 61 of the dialogue above. At this point, there remain a number of unexecuted actions in the plan. These remain on the goal stack of the system or student until their next turn to choose a goal. When their turn arrives, they will first attempt to re-satisfy the preconditions of these actions before choosing a new goal. Clearly, the student would be given

the option of returning to such unexecuted subgoals, or choosing a new plan. This would correspond to attempting to return to the dialogue state which existed before the interruption. Of course, it may be the case that the dialogue state has changed sufficiently - perhaps focus has shifted - so that these preconditions cannot be resatisfied, in which case a new plan is generated and conversed.

If we imagine that a the system operated with a much larger set of available dialogue goals, a question would remain as to whether a more coherent dialogue structure would be generated by pursuing disjuctive dialogue subgoals depth-first left-to-right, or with some other search algorithm. A futher problem to be addressed concerns when the system should decide to abandon its initial top-level concept/style plan and replan, or whether to attempt to do a 'local fixup', should the system fail to receive acceptance for goal pursuit. There are a number of possibilities, with various implications for computational effort to be expended and the degree to which the system should attempt to coerce the student. Clearly, such a question lies at the heart of the overall educational philosophy which underlies tutoring system design, but for the present we shall restrict ourselves to stating a number of logical possibilities:

(1) the system pursues its initial plan regardless of the student's conversational moves and non-acceptance of its dialogue goals;

(2) the system accepts any divergence from its plan on the part of the student, and completely replans;

(3) the system suggests a dialogue goal for the student to follow, but when the suggestion is not taken up (after concept c2/style s3) it preserves the initial plan

c1/s2, c2/s3, c3/s1, c4/s5

and tries to find some subplan which will produce a coherent plan from c2/s3 to c3/s1.

(4) when the student suggests some topic and goal which does not accord with the next step of its plan, it looks for some later position in its plan where this goal could be accomodated, and requests that the student save this topic until that point in the dialogue. If the student refuses, it replans.

There are other possibilites, but we shall not discuss them prematurely at this stage. In general,

we would envisage a situation where the system initially adopts the approach of (3) above, with respect to its plan, when it has evidence that the student has little knowledge with which to be able to choose appropriate learning goals, and that it gracefully withdraws to a position where this situation is reversed once it is sure that the student does possess such knowledge. This would correspond to something like position (2), where the system formed an initial plan but was willing to abandon it upon non-cooperation, since it would be confident that the student would be pursuing a viable learning goal. In fact, it is possible that in such a situation the system would not make a new *plan* at all, but would simply perform *situated action* (Suchman 1985), by responding to each of the student's moves in a manner which best matched its overall educational goals and approaches.

A critique of the model in relation to existing work

Topic and goal-based approaches to dialogue

The model which has been described can be viewed as an attempt to integrate a number of existing theoretical models from the Intelligent Tutoring Systems literature with a number of computational models for generation and comprehension of dialogue from the general linguistics literature. Dialogue models can be divided generally into those which assign a predominant role to macroscopic units based on *topic* or *focus*, and those which assign that role to *intentional* or *goal-based* units. The former topic-based approaches include the work of Grosz (1977), which has been expanded by Reichman (1985). Grosz's approach depends on partitioning a semantic network into focus spaces. We have similarly adopted a semantic network as the basis for representation of explicit and implicit focus, but we believe that a cognitive model for memory recall, such as the spreading activation model (Quillan 1968, Anderson 1984, McClelland & Rumelhart 1981), will provide a more cognitively realistic model for a conversant's conception of dialogue focus.

Intention based approaches view dialogue as fundamentally a mechanism which conversants use to achieve their goals, and that these goals can be achieved by an *integration* of physical and dialogue actions. Although a number of writers (Levin & Moore 1977, for example) have

adopted this approach, Power's (1979) model remains the only computational model to date. His approach was to develop a computational model for a conversation between two hypothetical robots, who generated dialogue 'actions' in order to negotiate cooperation to achieve joint goals. In the 'toy' world considered, such joint goals might be "robot John inside", and discourse actions were fully integrated with (hypothetical) physical actions to achieve a goal. As will be discussed, aspects of the model which we have described in the previous section are closely related to Power's overall approach of symmetry between dialogue participants, in goal based dialogues. In essence, our model takes his general approach and applies it within the context of an ITS, complete with domain and student models, and mechanisms for achieving educational goals.

Integrating the two approaches

The question as to which of these two approaches is 'correct' is probably not a fruitful one to attempt to answer here: clearly, any adequate model must attempt to *combine* intentional and topic structures, and the important issues concern the nature of this combination and the degree of precedence afforded to each component. Current research in computational models of dialogue (for example, Grosz & Sidner 1986) is in fact moving towards attempting to integrate these two complementary modes of analysis. From the ITS literature, the dialogue model which forms part of Elsom-Cook's (1984) IMPART system integrated topic structures with goal-based structures by essentially using goal based "dialogue games" (Levin & Moore 1977) to discuss possible topics identified during the user's interaction with a LISP environment. From the more general computational linguistics literature, Grosz and Sidner's (1986) work was an attempt to integrate Grosz's earlier (1977) work on dialogue focus with Sidner's and others' work on intention in dialogue. Fundamentally, our model can be viewed as a specific instance of their general thesis that

> "... the structure of any discourse is a composite or three distinct but interacting components:
>
> • the structure of the actual sequence of utterances in the discourse;
>
> • a structure of intentions;
>
> • an attentional state"

(p. 176).

Linguistic structure is concerned with the existence of linguistic "cues" to other kinds of structure (cf. Reichman's "clue words", 1985) , intentional structure expresses the purposes of units for discourse participants, and the attentional state corresponds to focus-based structures, modelled in terms of a set of focus spaces. A feature of their model which the model described in the previous section lacks, is a sophisticated conception of a participant's conception of various levels of discourse focus, in terms of a *focus stack*. It would be possible to implement this notion as part of our model by hierarchically *partitioning* (Hendrix 1977) the semantic network knowledge base.

Our approach to integrating views dialogue goal structures as the fundamental units of dialogue, with changes in topic viewed as *effects* of dialogue actions, constrained by a model of relevant focus shifts. This view concerning the primacy of intentional structures was adopted by Levin & Moore (1977). In their model, dialogue is viewed as the pursuit of a number of possible 'dialogue games', each of which had a specific *function*, or aims to achieve a specific conversational goal:

> "Changes of "topic" in dialogue are directly dependent upon changes in the participants' goal structures, and are thus accomplished as side effects of goal structure changes." (1977, p.400).

The specification of preannouncements designed to satisfy preconditions, and dialogue goal actions, can be viewed as 'dialogue games' with specific topics and speaker roles. For example, the dialogue game HELPING may be preceded by an INFO-PROBE dialogue game, as a presequence to discover whether the desired helper has the requisite knowledge which the helpee wants. We would not, however, make the strong claim that this primacy of intentional over topical structure would apply to all possible human dialogues. For example, Grosz (1977) has shown that for task-based dialogues dialogue structure follows the topical structure of the task domain. Our main concern is to use the model which has been described within educational contexts, where the importance of intentional structures is clear from the necessity to consider the teacher's educational goals and the pupil's learning goals.

Symmetry, cooperativity and negotiation

Important features of the dialogue model concern symmetry of goals available to system and student, and the necessity of ensuring cooperation on goal pursuit by negotiation. Clearly, these factors are connected, since there is little point in providing symmetrical dialogue structures if the system dispenses with negotiation and simply coerces the student to respond in the manner which satisfies its own goals.

Existing ITS's lack such a symmetrical model of dialogue, in that the user is generally treated as a passive participant in the interaction, who can - like a half empty bucket of impure knowledge - have facts poured in and out, and generally have impurities in knowledge filtered out by appropriate questioning. We would maintain that a model of the interaction should give more autonomy to the student as an active participant in the dialogue, and that models of the teaching interaction should move towards providing dialogue mechanisms which allow an appropriate cooperative balance to be achieved between the student's pursuit of his learning goals and the tutoring system's pursuit of tutorial goals.

Such a relationship between symmetrical dialogue models and negotiation was embodied in Power's model in a number of ways, which are shared with our model. Just as in Power's model, each conversing robot keeps a record of facts about the world and a model of the plan of action upon which they propose to cooperate, in the model we have described there is an analogous symmetry between

- the dialogue actions available to each conversant,
- the provision of symmetrical belief sets, and
- the creation of a plan for each participant.

In our model, each conversant has the same available set of dialogue goals, a plan is recorded and conversed for each, and each can adopt beliefs imparted by the other. In summary, dialogue symmetry is required to provide a cognitively realistic and intelligible dialogue model, and to allow the student to pursue learning goals in a manner required by the Guided

Discovery approach to tutoring.

In addition to the notion of symmetry in Power's (1979) model, a further important notion is that of securing *cooperativity* (Grice 1975) in pursuit of dialogue goals. As in Power's model, this is achieved by *explicit negotiation*, using *adjacency pairs* with a specific function, called *presequences*. The dialogue consists of a succession of intentional units preceded by a phase of metadiscourse concerning negotiation about the form of the future discourse, in addition to the structure of dialogue foci which crosscuts these other structures, and a larger level structure of tutorial styles. Negotiation is required in a tutorial context in order to ensure that the student has an opportunity to pursue learning goals, independently of a naive assumption that learning goals are the simple inverse of the tutor's educational goals. A problem which faces our (and Power's) approach of representing this negotiation phase explicitly, is that in real human to human dialogues much of the negotiation of speaker roles, turn taking, and so on, is done *implicitly*. Non-verbal human communication further enables dialogue to be an essentially *real-time* activity, in a way that present human computer dialogues occurring in computational time can not be. For example, in a natural spoken dialogue, the precise duration of pauses between utterances can be of considerable and varied communicational importance in specific contexts. Since we presently know of no means by which this essential temporal element could be incorporated into human-computer typewritten dialogues, we have simply adopted the convention that choice of dialogue goals alternates between speakers, but that each speaker has the possibility to refuse to cooperate on any specific goal. A further artificiality is that in natural language, negotiation of preconditions would usually be combined into a single utterance rather than several. We have adopted this simplification since our primary research focus is not the generation of natural language utterances.

Symmetry and belief revision

A further question to be considered in connection with symmetry between the user's and system's belief sets, is how to model a conversant's criteria for adopting an uncertain belief or not. In cases where a belief contradicts some generally held principle, then then a conversant may decide to not adopt a belief; but our problem is to find some cognitively realistic model for

belief acceptance on the part of the ITS which is perspicuous and intelligible to the student. A number of approaches exist to reasoning with uncertainty, which could be adopted on the part of the ITS, to decide when there was sufficient evidence to warrant adopting a belief. As Cohen has described (1985), these methods divide into those which adopt some probabilistic method, control methods (eg blackboard architectures), and 'diversification' methods, which attempt to use a set of criteria for deciding which of a set of likely candidate beliefs to adopt. We have little evidence that any of these methods are those which humans use, and so we have not attempted to incorporate any of these approaches into the dialogue model. However, it is clear that humans do make such decisions in some way, and that to preserve the *dialogue symmetry* which assumes that the student will adopt some of the beliefs imparted by the system, we need some kind of decision mechanism for determining when the system should adopt a belief put forward by the student. For the present we shall adopt the simplistic method of adopting a belief in preference to another when there are a greater *number* of justifying factors. In the context of justifying the existence of a musical grouping boundary, this criterion relates directly to Lerdahl & Jackendoff's (1984) 'grouping preference rule' (GPR 4) that a possible boundary which has a larger number of indicating factors is 'intensified'.

In addition to possessing an *uncertain* set of beliefs about musical structures, the knowledge base of the ITS also does not assume that it possesses all the possible candidate structures, and is to that extent *incomplete*. This follows from its willingness to adopt further possible beliefs advanced by the student. One future reseach possibility is to incorporate the cognitive model for *plausible* reasoning with incomplete knowledge, which has been developed by Collins & Michalski (1987), so that the system can develop cognitively valid methods for reasoning as to whether a belief advanced by the student is sufficiently plausible to be adopted.

The succession of dialogue units

A phase of negotiation in dialogue not only ensures cooperativity, but also has implications for the functional relationships between successive intentional and topical units in dialogue. The negotiation phase enables dialogues to be generated in which *any dialogue goal can follow*

any other, provided that it is agreed as a cooperative goal by means of negotiation. We would maintain that this represents a better model of human-like dialogue than certain *discourse genre* and *grammatical approaches*. Discourse genre approaches attempt to specify various dialogue genres - such as 'explanation' , 'argument', and so on - and, in effect, to specify normative schematic structures for these genres. For example, in the context of producing an expert system which can 'argue' for the acceptability of its knowledge base, Stutt (1987) has adopted the schematic argument structure described by Toulmin *et al* (1979). For the purposes of altering the knowledge base of an expert system, this may well be a suitable approach, but is unlikely to be able to do justice to the rich variety of naturally occuring tutorial dialogues. Certain approaches in tutorial dialogue which support a single tutorial style - such as 'Socratic tutoring' - can also be viewed as an attempt to specify a particular dialogue genre. It is axiomatic to our approach to developing a general dialogue model here that the model would have to support a number of such 'genres' in order to support different invididual tutorial and learning styles. A casual examination of naturally occurring dialogues reveals that conversants are able to integrate explanation, critique, questioning and various kinds of metadiscourse, freely within the same dialogue: no single genre called 'tutorial dialogue' can, we believe, be isolated. The fact that a single utterance can serve many different purposes in different contexts suggests that the goal which an utterance is intended to achieve is more fundamental than its membership of any single discourse genre. Since we wish to integrate critiquing claims, making claims, explaining fundamental principles, helping the student to form the best supporting arguments which they could produce, and so on, we have not adopted this approach.

Another possible candidate for relationships between dialogue units concerns the notion of *grammaticality*. For example, Reichman's (1985) dialogue model uses an ATN grammar to specify possible 'well-formed' transitions between topic-based units ("context spaces") using a set of available "conversational moves". A simple objection to any grammatical approach which attempts to apply formalisms which may operate at the sentence level to the dialogue level, is that *any* apparently 'ill-formed' segment of dialogue can appear completely intelligible, given a suitable context.

Conversants' expectations and adjacency pairs

The fact that, generally speaking,dialogues between people educated in certain cultural traditions exhibit certain regularities, does not reflect some universal grammaticality in discourse, or use of a finite set of discourse genres, but rather that given certain types of utterances, there are mutual *expectations* about what kinds of utterances may follow. These expectations may turn out to be unfulfilled and produce perfectly intelligible and natural sounding dialogues. One way in which these expectations can be characterised is in the dialogue phenomenon known as *"adjacency pairs"* (Schegloff & Sacks 1973), which we mentioned in discussion of preconditions for dialogue goals . For example,
the 'question-answer' pair

Michael: "What are you having?" **((question))**

Mark: "Mine's a pint of Hook" **((answer))**

One use of adjacency pairs which we have incorporated into our model is that of presequences, such as a preannouncement to a question-answer pair

Michael: "Hey, Mark" **((summons))**

Mark: "Yeah?" **((answer))**

Another use of such structures concerns the range of possible seconds to a first of an adjacency pair. Adjacency pairs to not in fact have to be adjacent, as for example in the following nested pairs, where the first response to the question-answer pair is in fact the first of another pair:

Michael: "Are you staying for another one?" **((question 1))**

Mark: "Have you got enough money on you?" **((question 2))**

Michael: "I think so" **((answer 2))**

Mark: "Alright, I'll have a Hook" **((answer 1))**

It is clear that two possible seconds are to give the expected *acceptance* response, or a dispreferred *rejection*. These are incorporated into our model by simply providing the alternative responses 'OK' or 'NOT OK', to the respondent. However, another way of *rejecting* to play the 'dialogue game' (Levin & Moore 1977) which the adjacency pair attempts to initiate (rather than flat rejection) is to reply with the first of a quite different adjacency pair, as in the

example above. So doing can represent an attempt by the respondent to pursue a different kind of game or goal. This is incorporated implicitly into our model by *alternation of goal choice:* after a respondent responds with 'NOT OK' to an initiator's attempted goal pursuit (the first of an adjacency pair), the goal is abandoned as failed (negotiation fails to secure an agreed goal), and the respondent pursues a goal using the first of a different adjacency pair in an attempt to secure negotiation preconditions of their goals. In this way, the first of one adjacency pair can receive the first of another as its (dispreferred) second. A feature which we have not been able to incorporate into the model is this *preference ordering* between (a larger number of) possible seconds, since we presently know of no theoretically interesting way of characterising context dependent ordering.

Tutorial strategies, dialogue goals and related work

Using the set of tutorial strategies and dialogue goals contained within the model which we have described, it would be possible to generate dialogues which are similar to a number of existing approaches. For example, a number of the dialogue goals which our model incorporates are similar to a number of Reichman's (1985) conversational moves. For example, our description of pursuit of goals remaining unpursued after an interruption , would be a 'return' move in terms of her theory, and pursuit of our 'complementary claim' goal would be an 'indirect challenge' move. There are also close affinities between the subgoals to 'challenge' goals which we have specified and approaches to *critiquing* (Miller 1984) and rules for *Socratic dialogue* (Collins 1977). There are, however, important differences between these approaches and the way in which certain elements of them have been incorporated into our model as dialogue goals. Critiquing approaches were developed with a view to improving user acceptance of expert systems in medical domains, by providing the user with a critique of their proposed treatment approach in terms of known risks of drugs and so on, rather than attempting to provide a definitive 'expert' answer. A critique is a single piece of generated text in response to a single proposed solution input by the user. The structure of the interaction is thus a simple answer/response pair, and as such cannot be considered to be a dialogue model in the sense in which we understand it. The model lacks a conception of goal and

topic structures, which is hardly surprising since it was never specifically intended to support educational interactions (although Miller hypothesises that there may be possible educational applications of the model). Similar points could be made about Collins' (eg 1977) approach to Socratic dialogues. In his 1977 work, the approach was to specify a set of *rules* for the tutor's dialogue action in response to certain antecedent conditions present in the student's previous response and certain features of the dialogue state. For example, his Socratic tutoring rule 6 is:

"Pick a counterexample for an insufficient factor.

if

(1) the student gives as an explanation one or more factors that are not sufficient, or

(2) agrees to the general rule in rule 5

then

(3) pick a counterexample that has the right value of the factor(s) given but the wrong value of the dependent variable, and

(4) ask what the value of the dependent variable is for that case, or

(5) ask why the causal dependence does not hold for that case."

In a sense, this is similar to Reichman's 'direct challenge' conversational move and to the complementary claim goal which we described. Although these rules do embody many certainly true facts (the rules were based on extensive protocols) about tutorial arguments which seek to develop students' critical faculties, we would question the rule-based framework as a formalism for quite low-level dialogue actions. It is possible for a number of tutoring rules to possibly apply at any given point in the interaction, and we are given no educationally based or non-arbitrary means of conflict resolution. In addition, the model lacks any larger scale topic and intentional structures which are macroscopically organised to give an intelligible dialogue structure. In a sense, it could be argued that at a sufficent level of abstraction there is no essential difference between having preconditons for an action which achieves a goal, and having antecedent conditions for executing the right-hand side of a rule. But the essential difference here is that the plan based formalism enables us to combine such actions into a sequence of actions which constitute a coherent dialogue, and one which can be used to achieve goals, in a manner which is not intuitively obvious with chains of rules (Socratic rules are applied singly, with no rational relationship between successive rule firings). Collins' later

work (1983) is however, moving more towards a recognition that the tutor's and student's goals need to be accomodated within the tutoring rule system, in terms of decision criteria for selecting cases to be discussed.

In summary, therefore, the principal differences between work on Socratic dialogue, critiquing and Reichman's conversational moves, and the dialogue model which we have discussed are that our model

- integrates intentional and focus based dialogue structures
- it has a conception of relevant dialogue focus shifts
- it incorporates features which give greater autonomy to the student in terms of symmetrical dialogue structures, and cooperativity through negotiation
- it integrates a number of applicable tutorial styles within a single dialogue model.

Conclusions

In conclusion we shall attempt to assess limitations and possible extensions of the model, and to clarify research contributions to computational generation of tutorial dialogues.

All existing ITSs (at least those described in the recent survey by Wenger 1987) make the assumption that the domain knowledge which they aim to impart is possessed as a collection of declarative facts which are assumed to be certainly true. Given a certain relationship between knowledge and its means of tutorial communication, this usually has the effect that tutorial strategies are developed which simply aim to coerce the student into acquiring that same body of knowledge. The research which has been described in this chapter thus questions two fundamental assumptions of most existing ITS research, which are that:

(1) an intelligent tutoring interaction is essentially characterised by *knowledge communication*;

(2) an intelligent tutoring interaction is essentially characterised by communication of *certain* knowledge.

To consider claim (1) first, the dialogue structure which we have described is not primarily concerned with communicating knowledge *per se*, but rather with the *use* of that knowledge. Our approach has thus something in common with any tutoring approach which concerns itself with a student's reasoning methods. Rather than being concerned with a form of reasoning which is *not* the way that humans reason - such as reasoning with first order predicate logics (Johnson-Laird 1983) - but which is viewed as some kind of ideal in a domain (such as mathematical reasoning by deduction), we have tried to address a form of reasoning which is closer to that which humans actually do, and which is a sufficiently general cognitive skill to be applicable to other subject domains. Other domains to which the dialogue model is directly applicable are therefore those in which uncertain knowledge and critical arguments predominate, such as archaeology (Stutt 1987), philosophy and ethics, politics, law, and any science which advances by means of critical argument. We have termed this form of reasoning by means of a naturalistic human dialogue, *critical argument*.

With respect to the second assumption (2) above, we would maintain that the assumption of *certain* knowledge in an ITS stems from a lack of concern with developing a cognitively realistic domain model. If we were to attempt to do this in many domains (such as mathematics), then rejecting the notion that the tutoring system should be a Machiavellian and omniscient mentor may lead to quite different tutorial styles being developed. No human teacher can solve every problem in a domain in which they have expertise, and if they could, then no teacher could do so exclusively with the most ideal solution methods. We would therefore maintain that the issue of developing appropriate tutorial styles which take account of uncertain knowledge (which we argued should be encompassed within a general dialogue model) is one which an ITS in *any* domain must address. The use of cognitively unrealistic domain models in certain domains partly stems from the existence of formalisations (such as in mathematics) . In our example domain of musical structures, there were no existing formal models, and so the development of such a domain model faced the opportunity to emphasise cognitive aspects in the creation of a new model.
It may be objected, that the model we have described does not in fact restrict itself exclusively

to knowledge which is assumed uncertain, since it assumes that it has represented the space of possible justifications, which are themselves assumed to be correct. The fact that this may be true does not mean that the *primary* set of beliefs which the system addresses are not assumed uncertain, and that the system does not therefore address the problem of tutoring with uncertain knowledge. It does do so, but within limitations which are common to all existing intelligent computer-based systems, which are that the knowledge base (of justifications) is assumed to be *closed*, and that we do not address the problem that the student may have a quite different conceptualisation of the domain from the system. Now clearly, no knowledge domain is 'closed' in this sense: *justification has no end*, since we can always ask for the justification of any justification advanced. We have adopted the simplification in the present model that the justifications which the system possesses are justified *a priori*, or assumed 'certain'.

An intriguing possibility for future research would be to make no such assumptions that justifications are themselves justified. Given a sufficiently wide vocabulary to describe existing features of a domain, which were intelligible to the ITS, we could conceive of the possibility of the system accepting a new factor as a possible justification, where a student could produce a consistent case that this was also present as a factor in instances of beliefs which the system already held. Given such a conception, then our system would be truly making *no* assumptions of certainty in its knowledge base: the system would be concerned purely with the way in which putative knowledge is utilised in human 'understanding'. A similar point of view on knowledge acquisition by a tutoring system has been recently proposed by Gilmour and Self (1988), who quote Michalski in claiming that

> "learning abilities are required not only from the student but from the tutor as well"
> (in Gilmour & Self 1988, p. 196)

We believe that such intelligent tutoring systems would embody a consistent, principled and egalitarian educational philosopy.

Future reseach would thus not be able to avoid a number of fundamental problems in ITS research. Those which we have already mentioned include:

- developing systems which can model a student's idiosyncratic

conceptualisation of a domain;

- developing systems with rationally extensible knowledge;
- developing cognitively valid models for when a belief should be accepted.

In addition, a number of more specific problems arise in the dialogue model. For example, we have not really addressed the problem of goal conflict between dialogue participants: we assumed that each participant will simply acquiesce gracefully when they do not receive complete acceptance for the preconditions of their proposed dialogue goal. Other problems include managing goal negotiation in less artificial manner, which relates directly to the narrow "communication channel" (Elsom-Cook 1984) through which an ITS can communicate. We have at least provided a perspicuous formal framework within which these research problems can be addressed.

Acknowledgements

Many thanks to my supervisor, Dr. Mark Elsom-Cook, has carefully guided, inspired, and contributed to all of the ideas presented in this chapter. Dr. Claire O'Malley and Dr. Ben du Boulay have also given much needed criticism. Thanks to Fiona Spensely for reading the manuscript, and to Genevieve and Mido for providing an enjoyable and creative working environment .

This research is supported by a postgraduate study award from the Science and Engineering Research Council of Great Britain.

References

Anderson, J. R. (1984). Spreading Activation. In Anderson, J. A. and Kosslyn, S. M. (eds.) Tutorials in Learning and Memory, W.H. Freeman, New York.

Baker, M.J. (1988a, in press). An Artificial Intelligence Approach to Musical Grouping Analysis. Forthcoming in Contemporary Music Review, special issue on Science and Music (c. October 1988).

Baker, M.J. (1989, in press). A Cognitive Model for Perception of Musical Grouping

Structures. Forthcoming in Contemporary Music Review, special issue on Music and the Cognitive Sciences (c. February 1989).

Bellack,A., Kliebard,H., Hyman,R. & Smith,F. (1966). The Language of the Classroom. Teachers College Press, Columbia University, New York.

Bonar, J.G. & Cunningham, R. (1988). BRIDGE: an intelligent tutor for thinking about programming. In Self, J. (ed.) Artificial Intelligence and Human Learning, Chapman Hall, London.

Clancey, W. (1987). GUIDON: Knowledge based tutoring, MIT Press.

Cohen, P.R. (1985) Heuristic Reasoning about Uncertainty: An Artificial Intelligence Approach , Pitman, London 1985.

Collins, A. (1977). Processes in acquiring knowledge. In Anderson, R.C., Spireo, R.J. and Montague, W.E. (eds.), Schooling and the Acquisition of Knowledge. Lawrence Erlbaum Associates, Hillsdale New Jersey.

Collins, A. & Stevens, A.L. (1983). A cognitive theory of inquiry teaching. In Reigeluth, C.M. (ed.) Instructional-design theories and models: an overview of their current status, Lawrence Erlbaum.

Collins, A., & Michalski,R. (1987). The logic of plausible reasoning: a core theory. Cognitive Science, in press.

Draper, S. (1987). Spontaneous Explanation: report on the dialogue section of the 3rd Alvey workshop on explanation. Proceedings of the 3rd Alvey Workshop on Explanation, University of Surrey (Great Britain), 29th September-1st October, 1987.

Elsom-Cook,M. (1984). Design Considerations of an Intelligent Tutoring System for Programming Languages. PhD thesis, Warwick University, Great Britain.

Fikes & Nilsson (1971). STRIPS : A new approach to the application of theorem proving to problem solving. Proceedings of the Second International Joint Conference on Artificial Intelligence, Imperial College, London September 1971.

Gilmour, D. & Self, J. (1988). The application of machine learning to intelligent tutoring systems. In Self, J. (ed.) Artificial Intelligence and Human Learning (1988), Chapman and Hall, London.

Grice, H.P (1975). Logic and Conversation. In Cole, P., & Morgan, J.L. (eds.) Syntax and Semantics 3: Speech Acts, 41-58. Academic Press, New York.

Grosz, B. J. (1977). The representation and use of focus in a system for dialogue understanding. Proceedings of the 5th International Joint Conference on Artificial Intelligence, M.I.T., Cambridge Massachusetts, USA.

Grosz, B.J. & Sidner, C.L. (1986). Attention, Intentions, and the Structure of Discourse. Computational Linguistics, 12, 3, 175-204.

Hendrix, G.G. (1977). Expanding the utility of semantic networks through Partitioning. Proceedings of the 4th International Joint Conference on Artificial Intelligence.

Johnson-Laird, P.N. (1983). Mental Models. Cambridge University Press, Cambridge, Great Britain.

Lerdahl, F. & Jackendoff, R. (1984) A generative theory of tonal music, MIT Press.

Levin, J.A. & Moore, J.A. (1977). Dialogue-games: metacommunication structures for natural language interaction. Cognitive Science vol. 1, part 4, pp. 395-420.

McClelland, J.L. & Rumelhart, D.E. (1981). An interactive model of context effects in letter perception: Part I. An account of basic findings. Psychological Review, 1981, 88, 375-407.

Miller, P. L. (1984) A Critiquing Approach to Expert Computer Advice: ATTENDING, Pitman Publishing, London.

Ohlsson, S. (1986). Some Principles of Intelligent Tutoring. Instructional Science, 14, no. 3-4, 293-326.

O'Malley, C. (1987). Applying studies of natural dialogues to human-computer interaction: are two heads better than one? Proceedings of the 2nd Alvey Workshop on Explanation, University of Surrey (Great Britain), January 8-9th, 1987.

Peachey, D.R.& McCalla, G.I. (1986). Using planning techniques in intelligent tutoring systems. International Journal of Man-Machine Studies, 24, 77-98.

Power, R. (1979). The Organisation of Purposeful Dialogues. Linguistics 17, (1979) , 107-152.

Quillan, M.R. (1968) . Semantic Memory. In Minsky, M. (Ed.), Semantic Information Processing, pp. 227-270, Cambridge Massachusetts, MIT Press.

Reichman, R. (1985). Getting Computers to Talk like You and Me. MIT Press, Cambridge Mass.

Schegloff, E.A. & Sacks, H. (1973). Opening up Closings. Semiotica, 7, 4, 289-327.

Shneiderman, B. (1982). The future of interactive systems and the emergence of direct manipulation. Behavior and Information Technology, 1, 237-256.

Spensley, F. & Elsom-Cook, M. (1988). Dominie: Teaching and assessment strategies. Open University CAL Research Group Technical Report No. 74, Milton Keynes, Great Britain.

Stevens, A.L., & Collins, A. (1977). The goal structure of a socratic tutor. BBN Technical Report no.3518.

Stutt, A. (1987). Second Generation Expert Systems . Explanations, Arguments and Archaelogy. Human Cognition Research Laboratory Technical Report number 25, The Open University, Milton Keynes.

Suchman, L.A. (1985). Plans and situated actions: the problem of human-machine communication. XEROX PARC technical report P85-00005, Palo Alto Research Center, Palo Alto.

Toulmin, S., Rieke, R., & Janik, A. (1979). An introduction to reasoning. Macmillan Publishing, New York.

Weiner, J.L. (1980). BLAH: A system which explains its reasoning. Artificial Intelligence**15**,(1980) , 19-48.

Wenger, E. (1987) .Artificial intelligence and tutoring systems: computational and cognitive approaches to the communication of knowledge. Morgan Kaufmann Inc., Los Altos, California.

Winston, P. (1977) Learning to identify toy block structures. In Johnson-Laird, P.N. & Wason, P.C. (eds.) Thinking : readings in cognitive science, Cambridge University Press.

CENTRE FOR INFORMATION TECHNLOGY IN EDUCATION

List of CITE Reports

These reports may be obtained from:

Hansa Solanki, Institute of Educational Technology, The Open University, Walton Hall, MILTON KEYNES, MK7 6AA, England.

Report No.	Title and Author
1	A.T. Vincent, (1985) Computing and the Blind.
2	A. Jones, G. Kirkup, J. Morrison (1985) A Trial of Home Based Computer Terminals.
3	Gill Kirkup, (1985) The Present and Potential Use of Ceefax in the Open University.
4	Mark Elsom-Cook, (1986) Artificial Intelligence and Computer Assisted Instruction.
5	Mark Elsom-Cook, (1986) A Pascal program checker.
6	Simon Holland, (1986) How computers are used in the teaching of music and speculations about how Artificial Intelligence could be applied to radically improve the learning of composition skills.
7	Simon Holland, (1986) Design consideration for a human-computer interface using 12-tone three-dimensional harmony space to aid novices to learn aspects of harmony and composition.
8	Alison Petrie-Brown, (1987) The Influence of Context and Coherence as a Foundation for Dialogue Research.
9	Eileen Scanlon, Randall B. Smith (1987) A Rational Reconstruction of a Bubble Chamber Simulation Using The Alternate Reality Kit.
10	Mark Elsom-Cook, (1987) Intelligent Computer-Aided Instruction research at the Open University.
11	Mark Elsom-Cook, (1987) Towards a framework for human-computer discourse.
12	Mark Elsom-Cook, (1987) MATILDA AND IMPART: Lisp tools.
13	Mark Elsom-Cook, (1987) Guided discovery tutoring and bounded user modelling in Intelligent Computer Aided Instruction ed. J. Self, Chapman-Hall 1987.
14	A.M. Petrie-Brown and M.T. Elsom-Cook, (1987) An Examination of an AI model of indirect speech acts.
15	A. Edwards, (1987) Integrating Synthetic Speech With Other Auditory Cues In Graphical Computer Programs For Blind Users.
16	S. Holland, (June 1987) A knowledge-based tutor for music composition.

17 S. Holland, (June 1987) New Cognitive Theories of Harmony Applied To Direct Manipulation Tools for Novices.

18 M. Baker. (July 1987) Intelligent Computer-Aided Instruction and Musical Performance Skills.

19 M. Baker, (August 1987) Proposed Research Directions for Intelligent Computer-aided Instruction in Musical Performance Skills.

20 A D N Edwards, (August 1987), Adapting interfaces for visually disabled users.

21 M Elsom-Cook, (September 1987), Acquisition of computing skills.

22 M. Baker, (September 1987), Computational Analysis of musical grouping structures.

23 M Baker, (June 1987), Automated Analysis of Musical Grouping Structures as a Basis for a Guided Discovery Environment for Interpretation of Music.

24 D Laurillard, (October 1987), The different forms of learning in psychology and education.

25 A.D.N. Edwards, (November 1987) Modelling blind users" interactions with auditory computer interface.

26 P. Fung, (November 1987) Novice Prolog Programmers.

27 P. Fung, B. DuBoulay & M. Elsom-Cook, (November 1987), An initial taxonomy of novices' misconceptions of the Prolog interpreter.

28 G. Kirkup, (November 1987), Considering the effect on women students of an increased use of microcomputers in distance education.

29 Sara Hennessy, Rick Evertsz, Dave Ellis, Phil Black, Tim O'Shea, Ann Floyd, Design Specification for 'Shopping on Mars' a computer-based Educational Activity.

30 A D N Edwards, The Use of home computers by disabled students at the Open University. Part 1: Previous use of computers in courses

31 Michael E. Fox (March 1988) A Report on Studies of Motivation, Teaching and Small Group Interaction with special reference to Computers and to the Teaching and Learning of Arithmetic.

32 Michael E. Fox (March 1988) Theory and Design for a Visual Calculator for Arithmetic.

33 M. Baker (February 1988) A Cognitive Model for Perception of Musical Grouping Structures

34 M. Baker (February 1988) An Artifical Intelligence Approach to Musical Grouping Analysis

35 P. Fung (December 1987) Novices' predictions of Prolog's control flow: A report on an empirical study

36 M. Elsom-Cook and F. Spensley (April 1988) Knowledge representation in a tutoring system for procedural skills.

37 F. Spensley & M. Elsom-Cook (April 1988) Dominie: Teaching and Assessment Strategies.

38 M. Baker (May 1988) Tutoring with Incomplete and Uncertain Knowledge

39 R. Moyse (May 1988) Multiple Viewpoints for Intelligent Tutoring Systems.

40 P. Whalley (May 1988) Hyper Technic - a graphic object-oriented control environment.

41 P. Whalley (June 1988) Cued recall as a Measure of argument integration.x

42 L. Alpay (June 1988) A Survey and Examination of Intelligent Tutoring Systems in Medicine

43 M. Elsom-Cook (July 1988) Introduction to the ECAL system.

44 F. Spensley (July 1988) Dominie: Trainer Interface.

45 M. Baker (July 1988) An Architecture of an Intelligent Tutoring System for Musical Structure and Interpretation.

46 E.Scanlon and R.Sibbitt (April 1988) Proceedings of CAL Conference 88.

47 L. Alpay (September 1988) Medical Problem-Solving and Intelligent Tutoring Systems: Proposed Research Directions.

48 S.A Cerri, M.T. Elsom-Cook, M. Leoncini (September 1988) TRILL: The Rather Intelligent Little Lisper.

49 R. Evertsz, S. Hennessy & R. Devi (September 1988) GADL: a Graphical Interface for Mental Arithmetic Algorithms.

50 P. Fung, (October 1988) A formalisation of novices' errors in Prolog programs.

51 M. Baker, (October 1988) Arguing with the tutor: a model for tutorial dialogue in uncertain knowledge domains.